9

David Papineau

Professor of Philosophy

King's College London, UK

How did you get interested in the philosophical aspects of the social sciences?

My first academic appointment was to a Lectureship in the Department of Sociology at the University of Reading in 1973. I suppose that this was a slightly odd appointment, given than my academic training until then had been entirely in mathematics and philosophy. But the Head of the Reading Department at that time was Stanislav Andreski, an émigré trained in the pre-war Polish logical empiricist tradition of comparative sociology. He thought that the undergraduates needed a solid grounding in basic logic and Mill's methods of induction and that I was the person to give it to them.

I was there for four years and it was a wonderful education for me. Anybody who has read Andreski's delightfully irreverent *Social Sciences as Sorcery* (1972) will know that he had no patience with sociological fashion. As a result, few of his appointments were conventional sociology PhDs. Instead he amassed a distinguished collection of intellectual waifs and strays. Many were from Europe, displaced by the war and its aftermath, and shunned by more conventional employers. The ethos of the department was empirical but by no means narrow. Bill Russell, who taught ethology and cultural evolution, had trained as a psychoanalyst. Viola Klein had been a student of Mannheim's and was one of the first sociologists to write on the role of women. Alexander Lopasic was an anthropologist from the then Yugoslavia, Maria Hercowitz was an industrial sociologist from Poland, and Peter (Tank) Waddington was a criminologist who had recently left the police force. Overall the focus of the Department was on the broad sweep of history, and it would be hard to imagine a collection of people who knew more about it.

My first book, *For Science in the Social Sciences* (1978), was based on my 'Methodology' lectures in Reading. In fact I didn't do much methodology in the lectures. Mostly I discussed the then current topics in the philosophy of social sciences: holism versus individualism, explanation versus understanding, historical determinism, traditional belief systems, fact and value. A few sections of the book now strike me as a bit dated, concerned with debates that have since gone dead, but for the most part I still stand by what I said then. As the title indicates, the book was an unabashed defence of the 'positivist' view that there is no great divide between the social and natural sciences. But I didn't say anything too extreme. Of course I denied that there is anything magical about humans or societies that puts them beyond the limits of normal empirical investigation. But beyond that most of the book consists of sensible attempts to unravel theoretical tangles and avoid seductive fallacies.

Curiously, the part of the book that I now like least is the section on general epistemology of science. My Head of Department had hired me to teach logic and Mill's methods, but mostly I'd been telling the undergraduates about Kuhn and Feyerabend. Not that I endorsed Kuhn and Feyerabend's relativism. I saw myself as upholding realism against their ingenious attacks on scientific objectivity. Except that it wasn't real realism that I was defending, but the sophisticated neo-Popperianism of Imre Lakatos. At that time philosophy of science in Britain had been hijacked by the Popperians, and we were all brought up to think that falsificationism was enough to make the world safe for science. Nobody pointed out that Popper and his followers weren't even close to realists, but simply out-and-out skeptics who insisted that it is always wrong to believe any scientific theory. I suppose it is understandable that I went along with this tosh—plenty of more distinguished figures than I were equally tainted—but I can't help feeling bad about it now.

Which social sciences do you consider particularly interesting or challenging from a philosophical point of view?

I'm disinclined to pick out any one particular social science as particularly philosophically challenging, if only because I find it difficult to take the institutional divisions between the different social sciences seriously in the first place. It has always seemed to me that the subject matters of sociology, anthropology, history, political science, social psychology and even economics are so intertwined as to make the traditional divisions between the disciplines seem artificial. I know that the different subjects have different institutional traditions, with different canonical figures and different ideas on how young scholars should be trained. But in my view the disciplinary introversion that this encourages is a vice, not a virtue. The best work in all areas is typically informed by thinking from other branches of social science.

One question worth asking about the social sciences taken collectively is—'what are they trying to find out?' Back in the 1970s I thought I knew the answer: the social sciences, like all empirical sciences, aim to uncover general laws. But now this seems to me far too narrow a conception of what the social sciences can achieve at a theoretical level. Not that I think that there are no serious laws in the social sciences. There are plenty of striking patterns supported by the empirical data, both at a large-scale and individual level. (If you are skeptical, here are two empirical hypotheses that have become prominent since the 1970s: there are no wars between democracies; there are no famines in democracies. Or, to switch to the individual level, consider the findings of Milgram's obedience experiment, or of ultimatum game experiments.) Still, it now seems to me quite misguided to hold, as I used to, that all interesting theoretical findings in the social sciences can be fitted into the mould of general law. . .

A first point here is one that applies to the natural as well as the social sciences. We can often know about the causes and mechanisms that operate in certain local contexts without knowing any laws that cover them. This is something that Nancy Cartwright has been stressing for years (1983, 1989, 2000). Causes typically come before laws, at least in the order of discovery. Cartwright originally stressed the point in connection with the physical sciences. But the moral applies all the more strongly in the social realm. Social investigations uncover huge numbers of causal truths, but very few laws.

Of course, we can agree with Cartwright on this methodological point without endorsing the heavy-duty metaphysical consequences that she takes to follow. I myself remain an unrepentant Humean about causes, at least to the extent of thinking that causes can be reduced to non-causal laws, and in particular to probabilistic laws that impose the requisite asymmetry on causal relationships (cf Papineau 1991). But I am happy to agree that, at a methodological level, the Humean constitution of causes is often epistemologically irrelevant. We typically pinpoint causes using techniques that don't require any identification of covering laws.

Randomized experiments make the point graphic. They can tell us that some treatment affects some result, but not how the size of this effect will vary with circumstances. The same point applies to 'Bayesian Net' methods for identifying causes from survey data. Indeed we can usefully view both these techniques as probabilistic analogues of even simpler forms of everyday reasoning codified as Mill's 'methods of induction'. To take the very simplest case, if some occurrence of interest is determined, and we know that only one change preceded it, then we can be sure that this was the operative cause in the circumstances, even if we don't know what other factors were required for it to produce the effect. Randomized experiments and survey analyses simply apply this everyday form of reasoning to probabilistic data.

As I said, these points about knowledge of causes and mechanisms apply to the natural as well as the social sciences. However, there is one specific kind of mechanism that does distinguish the social from the natural sciences. Many social processes are driven by mechanisms involving intentions. What happens depends on what people believe and desire.

I don't think that the structure of intentional mechanisms is nearly as well understood as people generally suppose. Back in the 1970s, along with pretty much everybody else, I was an unreconstructed internalist about intentional explanation. Intentions were internal states that explained actions. We used to argue about the status of these explanantions, and in particular about whether they yielded some distinctive kind of insight that was not on offer in the natural sciences. But even those who were 'anti-positivist' about such explanations still thought of them as explaining bodily actions, things you did in your own space, so to speak.

This internalist perspective now looks far too limited. Over the last few decades philosophers of mind have become persuaded that the contents of mental states are externally constituted. In principle, two thinkers could be molecule-for-molecule identical, yet their beliefs and desires refer to different things, because of differences in their environment and background. (Think of the same mental symbol having different meanings in the two cases, because it is being used to pick out different things—my physical duplicate is thinking of *his* wife, I'm thinking of *my* wife.)

Philosophers of psychology, especially Jerry Fodor, haven't liked this. We use intentional states to explain bodily behaviour, they urge, and externally constituted intentional states are no good for this. How you move depends on what is going on inside you. External differences that don't correspond to internal differences can't make a difference. (Fodor 1980, Segal 2000.)

I agree with Fodor that externally constituted states are no good for explaining bodily actions. But I don't think that this means that they are no good for explaining anything. The moral is that intentional explanations explain things in addition to bodily actions—namely, interactions between agents and their environments. The mistake is to suppose that intentional psychological explanation stops at the bodily peripheries. Rather, intentional explanation sees agency as extending out into the environment, and gives us a grip on how this works.

On reflection, it is scarcely surprising that intentional explanations should operate like this. After all, intentional states are individuated by truth or satisfaction conditions—that is, by specifications of how things have to be in the external world for beliefs to be true or desires to be satisfied. However, this representational dimension structure would surely be quite otiose if the only explanatory purpose of intentional states were to account internally for bodily behaviour. There's no obvious reason why the internal states that interact to cause bodily behaviour should also represent features of the external world. They would surely play their behaviour-generating role just as well even if they were 'blind syntactic' states whole nature was exhausted by their internal causal powers. (Cf Papineau 1993 ch 3.)

I would say that the explanatory point of ascribing intentional states is precisely to chart the impact agents will have on their environments when things go as they plan. From this perspective, there's nothing puzzling about the fact that intentional contents are often externally constituted—that's just what we should expect if we use them to track the way agents interact with their environments.

I'm not sure how far this issue is appreciated by philosophers of social science. The case for externalism about mental states has largely been made by analytic philosophers of mind, and has rested on intuitions about possible cases—twin-earth water, thigh-bone arthritis, swampman (Putnam 1975, Burge 1979, Davidson 1987). By contrast, the philosophers have paid relatively little attention paid to the special explanatory resources offered by externally constituted contents. This is one reason why internalists like Fodor have been able to argue that content externalism is an artefact of everyday intuitions and signifies nothing of explanatory importance. More work needs to be done in understanding the way intentional contents provide a bridge which relates the activities of agents to the wider world around them.

There is a great deal of interest among contemporary philosophers of social science in *collective* agency: do some supra-individual wholes qualify as agents in their own right? This is an important issue, but I would say that a prior task is to understand the structure of agency per se. Whether we are dealing with individual or collective agents, what exactly is that gets explained we explain their activities in intentional terms? It seem to me that we will have a much better understanding of the intentional mechanisms that are peculiar to the social sciences when we can answer this question.

How do you conceive of the relation between the social sciences and the natural sciences?

As I said earlier, I don't think that there is any great divide between the social and the natural sciences. People and societies are both parts of the natural world, and should be studied as such. Of course, there are particular investigative methods that are appropriate to particular social phenomena, just as there are particular methods appropriate to geology, hydrodynamics and plate tectonics. But this just shows that special kinds of complexity call for special investigative techniques, not that these different systems are constituted out of different fundamental ingredients.

On the question of fundamental ontological constitution, I am not just a naturalist but more specifically a physicalist. I think that everything is physically constituted. To invoke Saul Kripke's illuminating metaphor, once God had created all the physical stuff, his work was done—he had already made the people, thoughts, football clubs and nation states, by putting all the molecules in the right places.

Not that I take physicalism to be an a priori thesis. There is nothing contradictory in the idea that reality contains various kinds of non-physical stuff. This was certainly the orthodox opinion among mainstream scientists until little more than a century ago. But I take it that twentieth-century science, and in particular modern physiological research, has given us strong reason to doubt the existence of any non-physical entities. (Cf Papineau 2002 Appendix.)

Still, even if everything is physical, including the social realm, it doesn't follow that there are any interesting *epistemological* relationships between physics and the social sciences. In particular, it certainly doesn't follow that any claims made about the social realm need to be shown somehow to follow from the laws of physics. Non-fundamental natural sciences again provide a model. I take it that the subject matter of plate tectonics is physically constituted, if anything is. Yet this doesn't mean that the principles of plate tectonics need to be legitimated via some reduction to physics. Rather, they owe their acceptance directly to seismological, geological and other empirical evidence. Similarly, even if people and societies are physically constituted, we can seek to understand their workings by direct empirical means, without worrying about their physical underpinnings.

Is this a practical or principled point? It might seem that, if a system is physically constituted, then *in principle* it must be possible to understand it in physical terms, even if in practice this proves impossibly complicated. Some dispute even this much. Jerry Fodor (1974) and others have argued that the categories of the special sciences are 'variably realized' at the physical level, and that this constitutes an in-principle barrier to any uniform reduction of special scientific patterns to physical ones. To repeat Fodor's celebrated example, what counts as *money* will be physically quite different in different societies, so there can't possibly be any uniform physical story to explain the principles of monetary economics.

I have never been comfortable with Fodor's strong anti-reductionism about the special sciences. If the world is at bottom physical, shouldn't we expect systems with different physical realizations to behave differently? Putting it the other way around, if we observe some regular special scientific pattern, doesn't that argue that the special categories in question must have some uniform physical basis? From this perspective, I would argue that Fodor's money example is spurious. Money might be physically different in different societies, but it is *psychologically* the same everywhere—people expect others to provide goods and services for money, and therefore desire it themselves. And if psychological categories are uniformly realized at the physical level, then this would restore the in-principle possibility of a physical reduction of economic principles.

This kind of argument for in-principle reducibility is controversial. Some argue that natural selective processes of various kinds can give rise to genuine variable realizability and so block even in-principle reducibility. The idea is that different physical mechanisms often play the same psychological or social role precisely because they have all been *selected* to play that role. This then renders it unsurprising that the relevant mechanisms should be physically disparate—natural selection doesn't care, so to speak, about which mechanism plays some functional role, as long as it does the job. (Macdonald 1992, Papineau 1993 ch 2, Block 1997.)

So maybe Fodor's in-principle anti-reductionism can be saved by an appeal to Darwinian processes (which would be ironic, given Fodor's well-known antipathy to Darwinian thinking of any kind). Still, if we do side again the reducibility of the special sciences, it is important not to do so for the wrong reasons. Fodor sometimes writes as if in-principle reducibility would impugn the autonomy of the social sciences and mean that sociology departments ought to be taken over by physicists. But this does not follow at all. In-principle reducibility is a metaphysical thesis with no practical consequences. Even if the psychological and social realms were in principle reducible to physics, this wouldn't make any difference at all to practising social scientists. They would still investigate their subject by direct empirical means, with no concern for physical realizations.

By way of analogy, nobody thinks that molecules are variably realized at the physical level and that there is therefore an in-principle barrier to the reduction of chemistry to physics. But we still have chemistry departments as well as physics departments, for the familiar reason that it is impracticable to derive knowledge of chemical systems from our knowledge of the basic physical principles governing their components. The same point applies to sociology departments. We will still need them in practice, even if it is principle possible to reduce social truths to physics.

Which topics in the philosophy of social science will, and which should, receive more attention than in the past?

If there is an epistemological connection between the social and natural sciences, it involves biology rather than physics. Humans are biological beings as well as physical ones. While our physical nature may be largely irrelevant to our social being, the same is not necessarily true of our biology.

Not that social scientists have a very good record of thinking about this connection. Ever since the modern synthesis of Darwin and Mendel, the application of biological ideas to social phenomena has been blighted by a tendency to dichotomize human characteristics into the natural or the nurtured, the instinctual or the cultural, the innate or the acquired.

Of course, some improvements have been made. E.O. Wilson's 'sociobiology' (1975) corrected the naivety of the mid-twentieth-century assumption that any animal 'instincts' would automatically be for the good of the group. And the 'Evolutionary Psychology' movement of the 1990s corrected the naivety of sociobiological assumptions about the relation between contemporary behaviour and the past evolution of cognitive mechanisms. (Barkow, Cosmides and Tooby eds. 1992.)

But even the Evolutionary Psychologists tend to work with a crude and unexamined notion of 'innate' cognitive mechanisms. Part of the problem here is conceptual. What exactly is it for something to be 'innate'? Those who trade in this notion find it very hard to explain what they mean. 'Causally quite independent of environmental factors' is clearly far too strong—everything depends inter alia on environmental factor, including legs and arms. 'Unalterable by interventions' is also too strong, for just the same reasons. Perhaps the best that can be done, at least in connection with cognitive traits, is to equate 'innate' with 'normally develops without any help from learning or other psychological processes' (though this does leave us with the non-trivial task of explaining what 'normal' development and 'psychological' processes are). (Cf Griffiths 2002, Samuels 2002.)

But behind this conceptual awkwardness lies a more substantial issue. Are there really any innate cognitive traits? I myself think that our genetic heritage plays a huge role in shaping cognitive and hence social structure, but at the same time I doubt that *any* cognitive structures are 'innate', even in the weak sense of normally developing without learning. In my view, it's learning all the way down, right back to the earliest stages of ontogeny—but learning that is strongly biased from the start by genes that have been selected because they make us especially good at learning certain things.

There seem to me principled reasons for adopting this view. I've always wondered how the complex sets of genes supposedly determining Language Acquisition Devices, or Theories of Mind, or Cheater Detection Mechanisms, could have become fixed in our ancestral populations. There are well-known barriers to the fixation of such pluralities of genes, arising from their individual non-advantageousness. I can remember once airing this worry to my then student (and now colleague) Matteo Mameli. 'But what if proto-versions of these cognitive traits were originally due to general learning mechanisms?' he asked. 'Then you would expect a quick accumulation of genes that made people better at learning these traits.' I burst out laughing at the thought that all the prize exhibits in the nativist list of cognitive traits might be congealed versions of capacities that were originally due to general learning mechanisms. But of course Matteo was quite right. His suggestion happily explains why there is no evolutionary barrier to the selection of the genes that make us so very good at acquiring language, understanding other minds, and doing the other things that come so easily to us (Papineau 2005).

Note how this picture does justice to the 'poverty of the stimulus argument'—it explains why humans learn certain things with very minimal informational input—without requiring that any cognitive traits are 'innate' in the sense of requiring *no* learning. Indeed the picture gives us reason to doubt that there is ever innateness in this strong sense. After all, why should evolution bother to substitute genes for *all* informational input, so to speak, if that input is freely available in nearly all environments? (Babies could in principle have been genetically designed to 'grow' the ability to discriminate phonemes even if they never heard any human speech; but what would be the point, given how few suffer such impoverished environments?)

These considerations indicate that most of the important cognitive traits underlying human social organization are a product of *both* culture and genes. We have evolved certain genes because they made us good at acquiring certain intellectual traits; and these traits now come naturally to us because we have those genes. Of course, these genes also make us naturally good at other activities apart from those they were originally selected to foster (typing, mathematics, reading novels) and many of these are of great social importance. At the same time, the continued dependence of 'genetically natural' traits on learning means that they are by no means guaranteed to emerge in all modern environments.

The interdependence of genes and culture thus leaves us with many important questions. How do genes and culture interact in the course of individual ontogeny? How did they interact in the evolutionary history of hominid society? What limits do their interaction place on the space of possible societies? It is not clear how to answer these questions. We need more data (historical , comparative, developmental, and experimental) and more theories. But the first step is to abandon the outmoded assumption that cognitive traits can usefully be divided into those that are innate and those that are acquired.

References

Andreski, S. 1972. *Social Sciences as Sorcery* London: Andre Deutsch

Barkow, J. Cosmides, L. and Tooby, J. eds. 1992. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture* New York: Oxford University Press

Block, N. (1997) 'Anti-Reductionism Slaps Back' in Tomberlin, J. ed. *Philosophical Perspectives* 11

Burge, T. 1979. 'Individualism and the Mental' *Midwest Studies in Philosophy* 4

Cartwright, N. 1983. *How the Laws of Physics Lie* Oxford: Oxford University Press

Cartwright, N. 1989. *Nature's Capacities and their Measurement* Oxford: Oxford University Press

Cartwright, N. 2000. *The Dappled World: A Study of the Boundaries of Science* Cambridge: Cambridge University Press UP, 2000

Davidson, D. 1987. 'Knowing One's Own Mind.' *Proceedings and Addresses of the American Philosophical Association* 60

Fodor, J. 1974. 'Special Sciences or: The Disunity of Science as a Working Hypothesis' *Synthese* 28

Fodor, J. 1980. 'Methodological Solipsism considered as a Research Strategy in Cognitive Science' *Behavioral and Brain Sciences* 3

Griffiths, P. 2002. 'What is Innateness?' *Monist* 85

Macdonald, G. 1992 'Reduction and Evolutionary Biology' in Charles, D. and Lennon, K. eds *Reduction, Explanation and Realism* Oxford: Oxford University Press

Papineau, D. 1978. *For Science in the Social Sciences* London: Macmillan

Papineau, D. 1991. 'Correlations and Causes: Review Article of Nancy Cartwright *Nature's Capacities and their Measurement' British Journal for the Philosophy of Science* 42

Papineau, D. 1993. *Philosophical Naturalism* Oxford: Blackwell

Papineau, D. 2002. *Thinking about Consciousness* Oxford: Oxford University Press

Papineau, D. 2005. 'The Cultural Origins of Cognituve Adaptations' in O'Hear, A. ed *Philosophy, Biology and Life* Cambridge: Cambridge University Press

Putnam, H. 1975. 'The meaning of 'meaning' in Gunderson, K. ed. *Language, Mind and Knowledge, Minnesota Studies in the Philosophy of Science VII* Minnesota: University of Minnesota Press

Samuels, R. 2002. 'Nativism in Cognitive Science' *Mind and Language* 17

Segal, G. 2000. *A Slim Book about Narrow Content* Cambridge, Mass: MIT Press

Wilson, E. 1975. *Sociobiology: The New Synthesis* Cambridge, Mass: Harvard University Press